

Implementing Cisco IP Routing (ROUTE) (642-902)

QUESTION 21

When an OSPF design is planned, which implementation can help a router not have memory resource issues? Select the best response.

- A. Have a backbone area (area 0) with 40 routers and use default routes to reach external destinations.
- B. Have a backbone area (area 0) with 4 routers and 30,000 external routes injected into OSPF.
- C. Have less OSPF areas to reduce the need for interarea route summarizations.
- D. Have multiple OSPF processes on each OSPF router. Example, router ospf 1, router ospf 2

Answer: A

Explanation:

Memory issues usually come up when too many external routes are injected in the OSPF domain. A backbone area with 40 routers and a default route to the outside world would have less memory issues compared with a backbone area with 4 routers and 33,000 external routes being injected into OSPF. Router memory could also be conserved by using a good OSPF design. Summarization at the area border routers and use of stub areas could further minimize the number of routes exchanged.

The total memory used by OSPF is the sum of the memory used in the routing table (show ip route summary) and the memory used in the LSDB. The following numbers are a "rule of thumb" estimate. Each entry in the routing table will consume between approximately 200 and 280 bytes plus 44 bytes per extra path. Each LSA will consume a 100 byte overhead plus the size of the actual LSA, possibly another 60 to 100 bytes (For router links, this depends on the number of interfaces on the router). These amounts should be added to memory already used by other processes and by the IOS itself.

If you really want to know the exact number, you can do a show memory with and without OSPF being turned on. The difference in the processor memory used would be the answer.

QUESTION 22

The maximum number of routers per OSPF area typically depends on which three factors? (Choose three.)

- A. the kind of OSPF areas being implemented
- B. the number of external LSAs in the network
- C. the number of DRs and BDRs in the areas
- D. the number of virtual links in the areas
- E. how well the areas can be summarized
- F. the use of LSA filters

Answer: ABE

QUESTION 23

You are troubleshooting an OSPF problem where external routes are not showing up in the OSPF database Which two options are valid checks that should be performed first to verify proper OSPF operation? (Choose two.)

- A. Are the ASBRs trying to redistribute the external routes into a totally stubby area?
- B. Are the ABRs configured with stubby areas?
- C. Is the subnets keyword being used with the redistribution command?
- D. Is backbone area (area 0) contiguous?
- E. Is the CPU utilization of the routers high?

Get Latest & Actual [642-902 Exam's Question and Answers from PassLeader.](#)

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

Answer: AC

Explanation:

A totally stubby stubby area cannot have an ASBR so it will discard this type of LSA (LSA Type 5) -> A is a valid check.

Each stubby area needs an ABR to communicate with other areas so it is normal -> B is not a valid check.

When pulling routes into OSPF, we need to use the keyword "subnets" so that subnets will be redistributed too. For example, if we redistribute these EIGRP routes into OSPF:

```
+ 10.0.0.0/8+ 10.10.0.0/16+ 10.10.1.0/24
```

without the keyword "subnets"

```
router ospf 1 redistribute eigrp 1
```

Then only 10.0.0.0/8 network will be redistributed because other routes are not classful routes, they are subnets. To redistribute subnets we must use the keyword "subnets" router ospf 1 redistribute eigrp 1 subnets

-> C is a valid check.

We don't need to care if area 0 is contiguous or not -> D is not a valid check. CPU utilization cannot be the cause for this problem -> E is not a valid check.

QUESTION 24

When verifying the OSPF link state database, which type of LSAs should you expect to see within the different OSPF area types? (Choose three.)

- A. All OSPF routers in stubby areas can have type 3 LSAs in their database.
- B. All OSPF routers in stubby areas can have type 7 LSAs in their database.
- C. All OSPF routers in totally stubby areas can have type 3 LSAs in their database.
- D. All OSPF routers in totally stubby areas can have type 7 LSAs in their database.
- E. All OSPF routers in NSSA areas can have type 3 LSAs in their database.
- F. All OSPF routers in NSSA areas can have type 7 LSAs in their database.

Answer: AEF

QUESTION 25

When verifying OSPF virtual link problems, which is an important item to check on the two transit OSPF routers? Select the best response.

- A. OSPF process ID
- B. OSPF router ID
- C. OSPF network type
- D. OSPF memory usage
- E. OSPF CPU utilization
- F. OSPF stub area configurations

Answer: B

Explanation:

The OSPF router IDs of the two transit OSPF routers are used to form the virtual link (with the area area-id virtual-link neighbor-router-id command) so it is an important item to check -> B is correct.

QUESTION 26

You are developing a verification plan for an upcoming OSPF implementation. Part of this plan is to verify the status of type 3 LSAs within the network. Which routers should you verify first to ensure that the configurations are correct for generating type 3 LSAs?

Select the best response.

Get Latest & Actual 642-902 Exam's Question and Answers from PassLeader.

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

- A. Internal routers within the backbone area (area 0)
- B. Internal routers within the NSSAs
- C. Internal routers within the stubby areas
- D. ASBRs
- E. ABRs
- F. DRs and BDRs

Answer: E

Explanation:

Type 3 LSA (Summary LSA) is advertised by the ABR of originating area to advertise network from other areas so we should check the ABRs first.

QUESTION 27

Which condition must be satisfied before an EIGRP neighbor can be considered a feasible successor? Select the best response.

- A. The neighbor's advertised distance must be less than or equal to the feasible distance of the current successor.
- B. The neighbor's advertised distance must be less than the feasible distance of the current successor.
- C. The neighbor's advertised distance must be greater than the feasible distance of the current successor.
- D. The neighbor's advertised distance must be equal to the feasible distance of the current successor.
- E. The neighbor's advertised distance must be greater than or equal to the feasible distance of the current successor.

Answer: B

Explanation:

The feasible successor route is a route which has a higher metric than the successor route to reach a subnet but meets the feasibility condition and can be used in the event that the successor route goes down. This route does NOT get installed in the routing table but is kept in the topology table. The feasibility condition states that the AD from a neighbor must be less than the metric of the successor route (the feasible distance [FD]) because routing through a feasible successor when the AD > FD may cause a routing loop.

QUESTION 28

Based on the need to limit processing and bandwidth utilization due to dynamic routing protocol operation, the following routing requirements have been specified for your network.

- partial and incremental routing updates
- only the devices affected by a topology change perform route recomputation
- route recomputation only occurs for routes that were affected

Which dynamic routing protocol should be deployed in your network to best meet these requirements? Select the best response.

- A. BGP
- B. OSPF
- C. IS-IS
- D. EIGRP
- E. RIPv2

Answer: D

Get Latest & Actual [642-902 Exam's Question and Answers from PassLeader.](#)

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

Explanation:

The bandwidth utilization issue has been addressed by implementing partial and incremental updates. Therefore, only when a topology change occurs does routing information get sent. Regarding processor utilization, the feasible successor technology greatly reduces the total processor utilization of an AS by requiring only the routers that were affected by a topology change to perform the route recomputation. Furthermore, the route recomputation only occurs for routes that were affected. Only those data structures are accessed and used. This greatly reduces search time in complex data structures.

QUESTION 29

Which statement about a non-zero value for the load metric (k2) for EIGRP is true? Select the best response.

- A. A change in the load on an interface will cause EIGRP to recalculate the routing metrics and send a corresponding update out to each of its neighbors.
- B. EIGRP calculates interface load as a 5-minute exponentially weighted average that is updated every 5 minutes.
- C. EIGRP considers the load of an interface only when sending an update for some other reason.
- D. A change in the load on an interface will cause EIGRP to recalculate and update the administrative distance for all routes learned on that interface.

Answer: C

Explanation:

The load metric (k2) represents the worst load on a link between source and destination. EIGRP routing updates are triggered only by a change in network topology (like links, interfaces go up/down, router added/removed), and not by change in interface load or reliability. The load is a five minute exponentially weighted average that is updated every five seconds (not five minutes) . EIGRP considers the load of an interface only when sending an update for some other reason (like a link failure)

QUESTION 30

Your network consists of a large hub-and-spoke Frame Relay network with a CIR of 56 kb/s for each spoke. Which statement about the selection of a dynamic protocol is true? Select the best response.

- A. EIGRP would be appropriate if LMI type ANSI is NOT used.
- B. EIGRP would be appropriate, because the Frame Relay spokes could be segmented into their own areas.
- C. EIGRP would be appropriate, because by default, queries are not propagated across the slow speed Frame Relay links.
- D. EIGRP would be appropriate, because you can manage how much bandwidth is consumed over the Frame Relay interface.

Answer: D

Explanation:

By default, EIGRP will limit itself to using no more than 50% of the interface bandwidth. The primary benefit of controlling EIGRP's bandwidth usage is to avoid losing EIGRP packets, which could occur when EIGRP generates data faster than the interface line can absorb it. This is of particular benefit on Frame Relay networks, where the access interface bandwidth and the PVC capacity may be very different.

QUESTION 31

Get Latest & Actual [642-902 Exam's Question and Answers from PassLeader.](#)

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

When an EIGRP topology change is detected, what is the correct order of events when there is a FS? Select the best response.

- A. The neighbor adjacency is deleted.
The feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
- B. DUAL is notified.
Remove all topology entries learned from that neighbor.
The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
- C. The neighbor adjacency is deleted.
Routes enter the Active state and the feasible route is used.
DUAL is notified.
Remove all topology entries learned from that neighbor.
- D. DUAL is notified.
The neighbor adjacency is deleted.
Remove all topology entries learned from that neighbor. The feasible route is used.

Answer: D

Explanation:

If a packet is not received before the expiration of the hold time, the neighbor adjacency is deleted, and all topology table entries learned from that neighbor are removed, as if the neighbor had sent an update stating that all the routes are unreachable. If the neighbor is a successor for any destination networks, those networks are removed from the routing table, and alternative paths, if available, are computed. This lets the routes quickly reconverge if an alternative feasible route is available.

QUESTION 32

Refer to the exhibit. Why is the 140.140.0.0 network not used as the gateway of last resort even though it is configured first? Select the best response.

```
R3#show run | include default-  
ip default-network 140.140.0.0  
ip default-network 130.130.0.0  
R3#show ip route | begin Gateway  
Gateway of last resort is 0.0.0.0 to network 130.130.0.0  
116.0.0.0/8 is variably subnetted, 5 subnets, 3 masks  
C 116.16.37.0/30 is directly connected, Serial1/0.2  
C 116.16.32.0/30 is directly connected, Serial2/0.2  
C 116.16.34.0/28 is directly connected, Serial1/0.1  
C 116.16.35.0/28 is directly connected, Serial2/0.1  
S 116.0.0.0/8 [1/0] via 116.16.34.0  
* 140.140.0.0/32 is subnetted, 3 subnets  
O 140.140.1.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1  
O 140.140.3.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1  
O 140.140.2.1 [110/65] via 116.16.34.4, 00:14:54, Serial1/0.1  
* 130.130.0.0/16 is variably subnetted, 4 subnets, 2 masks  
D* 130.130.0.0/16 is a summary, 00:30:04, Null0  
C 130.130.1.0/24 is directly connected, Ethernet0/0  
C 130.130.2.0/24 is directly connected, Ethernet0/1  
C 130.130.3.0/24 is directly connected, Ethernet1/0
```

Get Latest & Actual [642-902 Exam's Question and Answers from PassLeader.](#)

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

D 150.150.0.0/16 [90/679936] via 116.16.35.5, 00:02:58, Serial2/0.1

- A. The last default-network statement will always be preferred.
- B. A route to the 140.140.0.0 network does not exist in the routing table.
- C. Default-network selection will always prefer the statement with the lowest IP address.
- D. A router will load balance across multiple default-networks; repeatedly issuing the show ip route command would show the gateway of last resort changing between the two networks.

Answer: B

Explanation:

As you can see in the exhibit, 140.140.0.0 doesn't appear in the routing table.

QUESTION 33

Refer to the exhibit. Why are the EIGRP neighbors for this router not learning the routes redistributed from OSPF? Select the best response.

```
router eigrp 123 redistribute ospf 123
network 116.16.35.0 0.0.0.255 network 130.130.0.0
auto-summary
!
router ospf 123
log-adjacency-changes
network 116.16.34.0 0.0.0.255 area 0
neighbor 116.16.34.4
```

- A. Redistribution must be enabled mutually (in both directions) to work correctly.
- B. Auto-summary causes the OSPF routes redistributed into EIGRP to be summarized; thus the OSPF network 116.16.34 is summarized to 116.34.0.0, which is already covered by the EIGRP protocol.
- C. Default metrics are not configured under EIGRP.
- D. Both routing protocols must have unique autonomous system numbers for redistribution to function correctly.

Answer: C

Explanation:

Same as RIP, when redistribute into EIGRP from OSPF, the default metric is infinite -> We must set a seed metric when redistributing into EIGRP.

QUESTION 34

Which BGP option is required when load sharing over multiple equal-bandwidth parallel links from a single CE router to a single ISP router over eBGP? Select the best response.

- A. eBGP Multipath
- B. eBGP Multihop
- C. BGP Synchronization
- D. Public AS numbers

Answer: B

QUESTION 35

Which BGP feature should be used to avoid high memory utilization on a router? Select the best response.

Get Latest & Actual [642-902 Exam's Question and Answers from PassLeader.](#)

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

- A. soft-reconfiguration
- B. route refresh
- C. BGP communities
- D. full-mesh BGP peering

Answer: B

Explanation:

BGP routers have enormous routing tables so it uses much memory to proceed these routes. When a BGP policy is changed, the BGP session needs to be reset for the policy to take effect. But the resetting results in route churn and route flapping. There are two ways to clear a BGP session without resetting the TCP session between them (this is often called "soft reset"):

Soft-reconfiguration: stores all received (inbound) routing policy updates without modification in a table so that when a new filter is applied, the router will use this table to calculate the changes without resetting the TCP session between the two BGP peers. This is a memory-intensive (high memory utilization) method and is not recommended.

Route-refresh: allows a BGP router to request a remote peer resend its BGP Adj-RIB-Out. This allows the BGP router to reapply the inbound policy. The route-refresh capability requires no extra memory on the local router

QUESTION 36

Which functionality is required within an IP router that is situated at the boundary of an IPv4 network and an IPv6 network to allow communication between IPv6-only and IPv4-only nodes? Select the best response.

- A. Autoconfiguration
- B. Automatic 6to4 Tunnel
- C. Automatic 6to4 Relay
- D. Network Address Translator-Protocol Translator (NAT-PT)
- E. Intrasite Automatic Tunnel Address Protocol (ISATAP)

Answer: D

Explanation:

http://www.cisco.com/en/US/prod/collateral/iosswrel/ps8802/ps6969/ps1835/prod_white_paper09186a008011ff51.html (Introduction, see 4th para)

QUESTION 37

The administrator wants to verify the current state of the OSPF database loading process. Which show command should the administrator use? Select the best response.

- A. show ip ospf [process-id] interface
- B. show ip ospf neighbor
- C. show ip ospf [process-id]
- D. show ip ospf [process-id area-id] database

Answer: B

The "show ip ospf neighbor" command can be used to view the current state of the OSPF database loading process.

QUESTION 38

During the IPv6 autoconfiguration, what does the device append to the 64-bit prefix that it receives from the router to create its IPv6 address? Select the best response.

Get Latest & Actual 642-902 Exam's Question and Answers from PassLeader.

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

- A. a pseudorandom generated number
- B. its locally configured IPv4 address
- C. the DHCP-supplied device ID
- D. its MAC address

Answer: D

Explanation:

The automatic configuration is a great feature of IPv6. Imagine you have to manually configure an IPv6 address with 128-bit long, what a pain! With this feature, it is no longer necessary to configure each host manually. But notice that host only autonomously configures its own Link-local address (the IP address used on a LAN). The Link-local address can be created automatically using a link-local prefix of FE80::/10 and a 64-bit interface identifier (based on 48-bit MAC address).

For example, if your MAC address is 00:12:34:56:78:9a, your 64-bit interface identifier is 0012:34FF:FE56:789a (16-bit FFFE is inserted in the middle). And notice that the notation has been changed because IPv6 addresses require 16-bit pieces to be separated by ":". Then, according to the RFC 3513 we need to invert the Universal/Local bit ("U/L" bit) in the 6th position of the first octet (start counting from 0). The "u" bit is set to 1 to indicate Universal, and it is set to zero (0) to indicate local scope. In this case we set this bit to 1 because the MAC address is universally unique. Thus the result is: 0212:34FF:FE56:789a. Finally, add the link-local prefix FE80 to create the full IPv6 address:

FE80:0:0:0:0212:34FF:FE56:789a (or FE80::212:34FF:FE56:789a, in short form). Note: The reason for inverting the "U/L" bit is to allow ignoring it for short values in the manual configuration case. For example, you can manually assign the short address fc80::1 instead of the long fc80:0:0:0:0200::1.

QUESTION 39

Which two methods use IPsec to provide secure connectivity from the branch office to the headquarters office? (Choose two.)

- A. DMVPN
- B. MPLS VPN
- C. Virtual Tunnel Interface (VTI)
- D. SSL VPN
- D. PPPoE

Answer: AC

Explanation:

<http://www.ccnpguide.com/ccnp-route-642-902-vpns-and-ipsec/>

QUESTION 40

You have implemented mutual route redistribution between OSPF and EIGRP on a border router. When checking the routing table on one of the EIGRP routers within the EIGRP routing domain, you are seeing some, but not all of the expected routes.

What should you verify to troubleshoot this problem? Select the best response.

- A. The border router is using a proper seed metric for OSPF.
- B. The border router is using a proper seed metric for EIGRP.
- C. The administrative distance is set for OSPF and EIGRP.
- D. The missing OSPF routes are present in the routing table of the border router.

Get Latest & Actual [642-902 Exam's Question and Answers from PassLeader.](#)

Click Here -- <http://www.passleader.com/642-902.html>

Implementing Cisco IP Routing (ROUTE) (642-902)

E. The subnet keyword on the border router in the redistribute OSPF command.

Answer: D

Explanation:

We are checking the routing table on EIGRP routers not OSPF so we don't need to check the seed metric for OSPF. Besides OSPF doesn't need to specify seed metric as all external routes get a default metric of 20 (except for BGP, which is 1). We must specify seed metrics when redistributing into EIGRP (and RIP). If not all the redistributed routes will not be seen but the question says only some routes are missing. The default administrative distance for external routes redistributed into EIGRP is 170 so we don't need to set it .

The subnet keyword is only used when redistributing into OSPF, not to other routing protocols . We should check the routing table of the border router to see the missing OSPF routes are there or not. An incorrect distribute-list can block some routes and we can't see it in other EIGRP routers.